MACHINE LEARNING FOR PARTICLE ACCELERATOR CONTROL SYSTEMS

Auralee Edelen

Fermilab New Perspectives Meeting

8-9 June, 2015

Abstract

Particle accelerators are host to myriad nonlinear and complex physical phenomena. They often involve a multitude of interacting systems, are subject to tight performance demands, and should be able to run for extended periods of time with minimal interruptions. Machine learning constitutes a versatile set of techniques that are particularly well-suited to modeling, control, and diagnostic analysis of complex, nonlinear, and time-varying systems, as well as systems with large parameter spaces. Consequently, the use of machine learning- and mathematical optimization-based modeling and control techniques could be of significant benefit to particle accelerators. For the same reasons, particle accelerators are also extremely useful test-beds for these techniques. This talk briefly discusses some promising avenues for incorporating machine learning into particle accelerator control systems and shows some initial results from our work at Fermilab.

Control Challenges for Particle Accelerators

- Particle accelerators are host to myriad complex/nonlinear physical phenomena
- Often involve a multitude of separately-controlled, interacting systems
- Can have many un-modeled disturbances
- Instabilities, coupling
- Long-term process cycles and drift
- Sometimes have limited diagnostics (large number of variables to adjust and just a few measureable outputs)
- Increasingly tight performance demands (transition to applications, increased energy/intensity)
- Desirable to run for extended periods with minimal downtime

What capabilities do we want?

- Automatically distill large amounts of data into useful information
  - Even for cases where data analysis is not straightforward
- Account for un-modeled behavior
- Take pre-emptive control actions if need be
- Adapt to drift in system behavior
  - A change in the rules governing how actions are decided upon (a “policy”)
  - Adaptation of a model (in model-based control)
- Find the best control actions to achieve a desired set of outputs
  - Reference tracking
  - Minimizing/maximizing some parameter(s)
  - Strictly adhere to constraints

→ *Much of this is can be addressed with machine learning and optimization*

→ *Particle accelerators can benefit from (and operate as a test-bed for) machine learning- and optimization-based control/data processing*

→ *Application attempts can help to guide theoretical development*

A trip to the zoo…

Artificial Intelligence

Machine Learning

Regression
Classification
Clustering
Dimensionality reduction

Learning Theory

Supervised Learning
Unsupervised Learning
Reinforcement Learning

Mathematical Optimization

Gradient descent
Conjugate gradient
Newton method
Quasi-Newton methods

Simulated annealing
Evolutionary algorithms
Swarm intelligence

Reactive search optimization

Computational Statistics

Biological Sciences *(inspiration!)*

Intelligent Control

Nonlinear Control

Adaptive Control

Optimal Control

Robust Control

Online data analysis (e.g. for diagnostics)

Fuzzy Logic

Expert Systems

Model-independent methods

System Identification

Model-based methods

Manual control and tuning by an operator:

- *Start from a previously known state or a state that is predicted to be acceptable*
- *Make a change, wait, observe a result, make another change based on some update rule*
- *For an operator the update rule might be based on understanding of the physics plus some memory of the immediately previous outcomes and general previous experience with similar systems, or it might be error-based*
- *Minimize difference between desired outcome and observed outcome*
- *Assess when an undesirable machine state has been entered*

Analogies:

- Grey-box modeling (mix of analytic theory, numerical simulation data, and measured data)
- Model predictive control (prediction/optimization of actions over a future time horizon)
- Model adaptation via online learning
- Reinforcement learning (regulated exploration of parameter space + creation of stimulus-response rules (a policy) + environmental feedback)
- High-level interpretation of data: clustering, classification, dimensional reduction

Many failures in the early days → so what's different now?

In general:

greater theoretical understanding

+

increased computational capability

+

advantageous co-developments in related fields

+

feedback from a wider variety of relevant application attempts

(and numerous successes in complicated offline data analysis tasks, process control tasks, fault prevention tasks, etc.)

→ ***greater overall technological maturity***

Changing Gears:
High-level Overview of Preliminary Work At Fermilab

Resonance Control for an Electron Gun

- 1½ cell, normal-conducting RF photoinjector operating at 1.3 GHz
- Water-cooled
- 23-kHz shift in resonant frequency per °C change in cavity temperature
- Existing requirements state that the water temperature should be regulated to within ± 0.02 °C

Simplified Water System Schematic

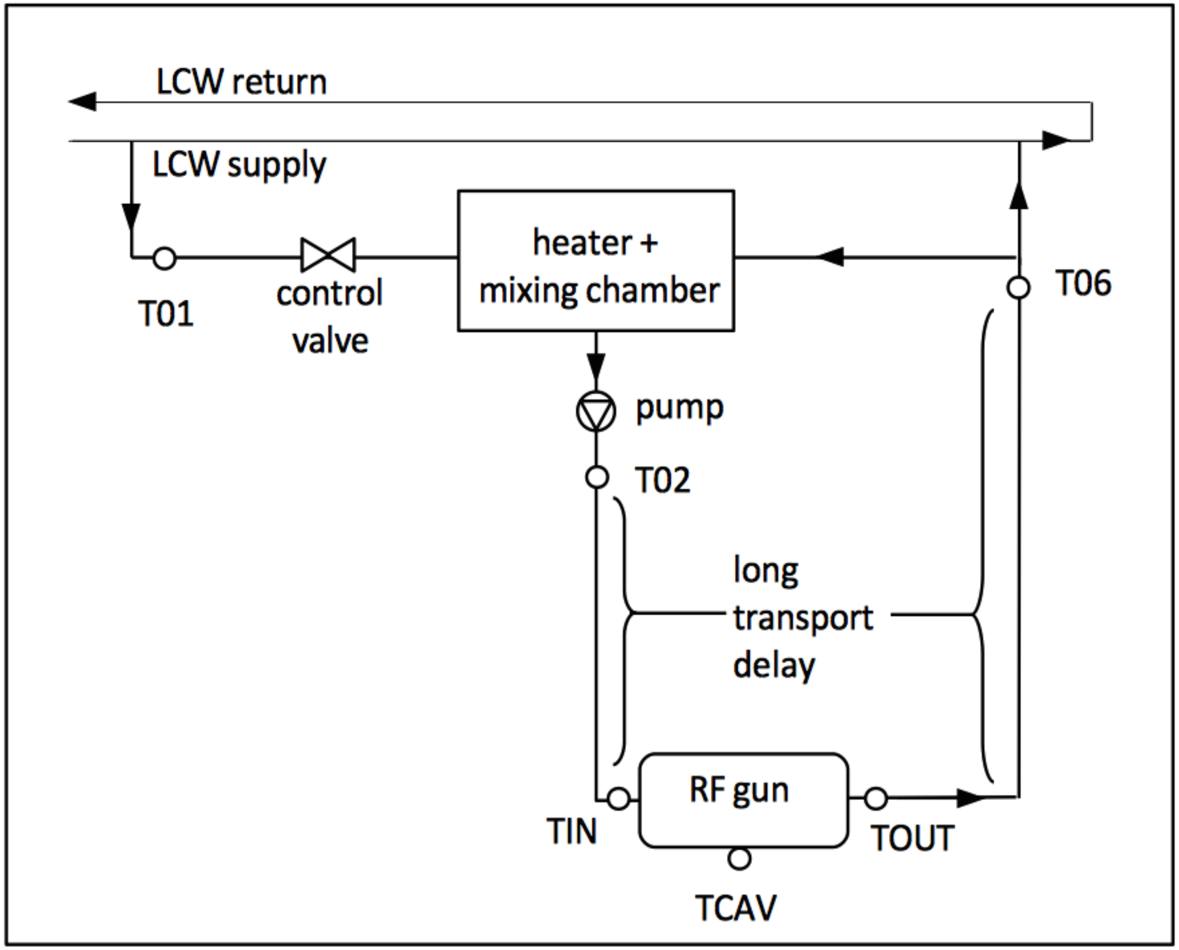

Division of Control Tasks

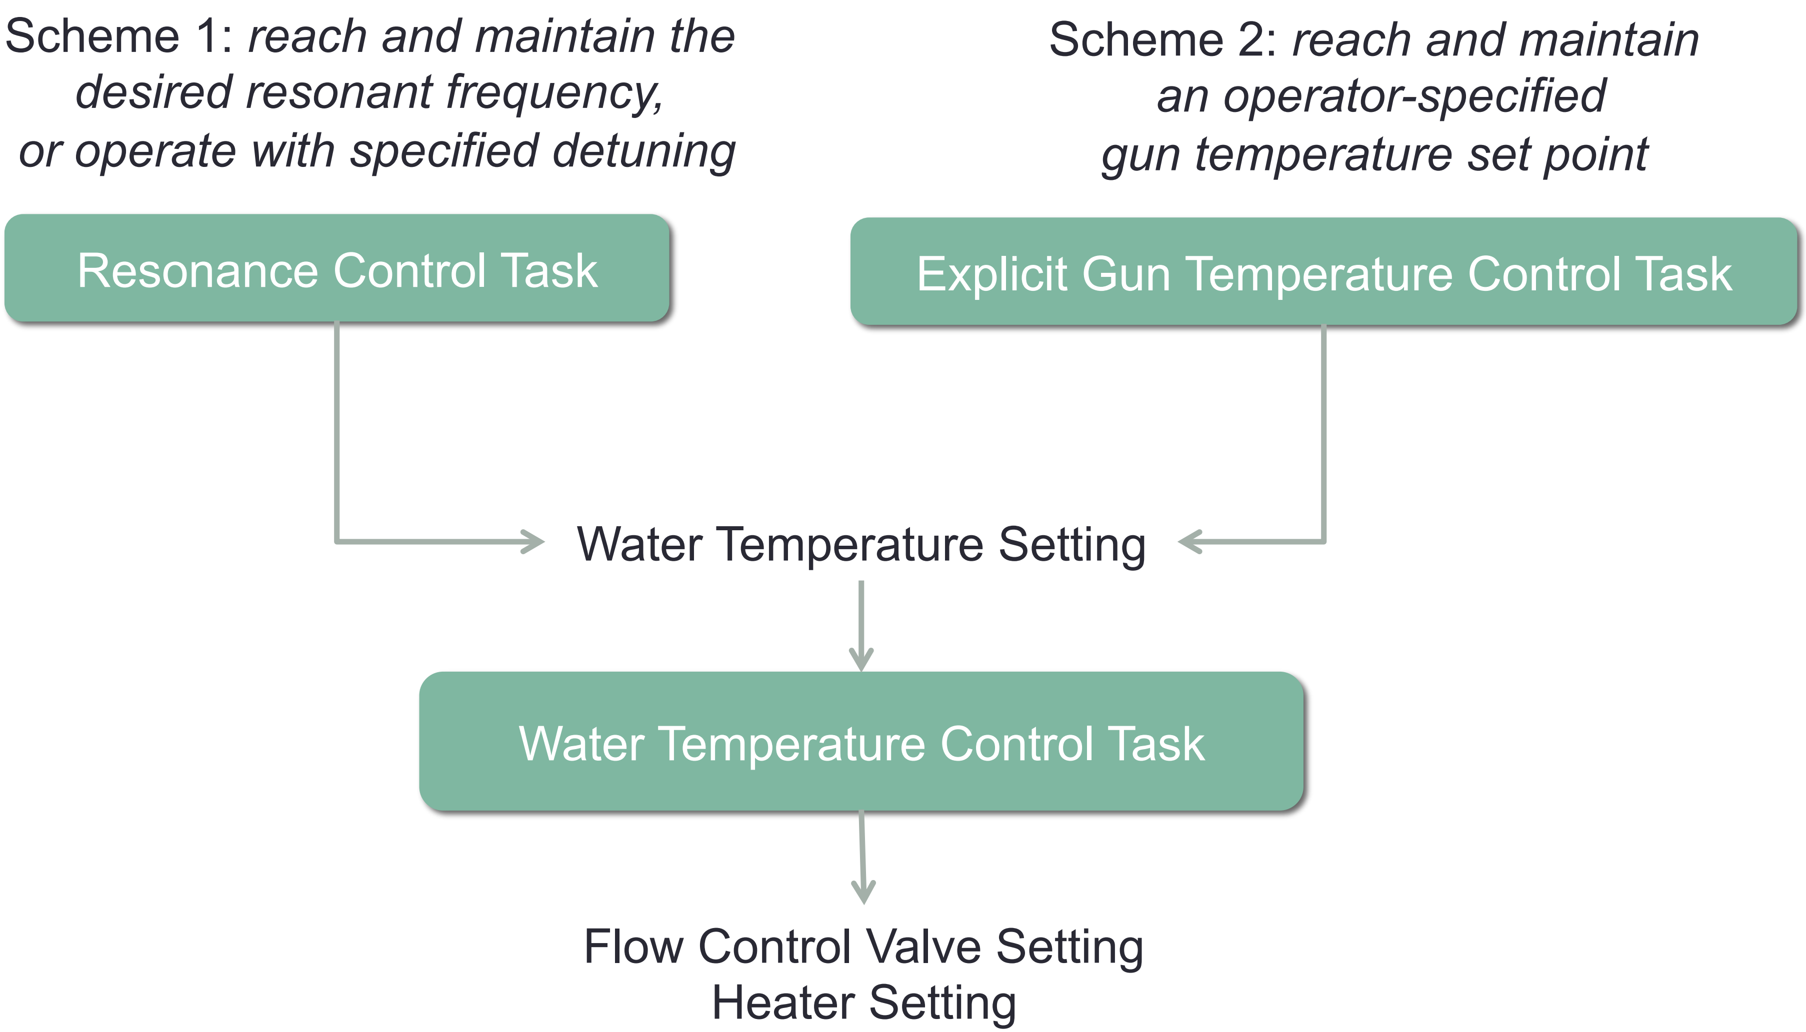

Division of Control Tasks

Scheme 1: *reach and maintain the desired resonant frequency, or operate with specified detuning*

Resonance Control Task

reinforcement learning

Scheme 2: *reach and maintain an operator-specified gun temperature set point*

Explicit Gun Temperature Control Task

simple identified/ adaptive model

Water Temperature Setting

Water Temperature Control Task

identified/adaptive model
+
predictive control (optimization)

Flow Control Valve Setting
Heater Setting

Division of Control Tasks

Scheme 1: *reach and maintain the desired resonant frequency, or operate with specified detuning*

Note: the actual cavity temperature is not necessarily well-represented directly by the temperature reading, as the region around the sensor experiences additional heating under RF power (linear relationship)

Can't take the readings strictly at face-value

Resonance Control Task

reinforcement learning

Water Temperature Setting

Water Temperature Control Task

identified/adaptive model
+
predictive control (optimization)

Flow Control Valve Setting
Heater Setting

Performance of Existing PID Controller

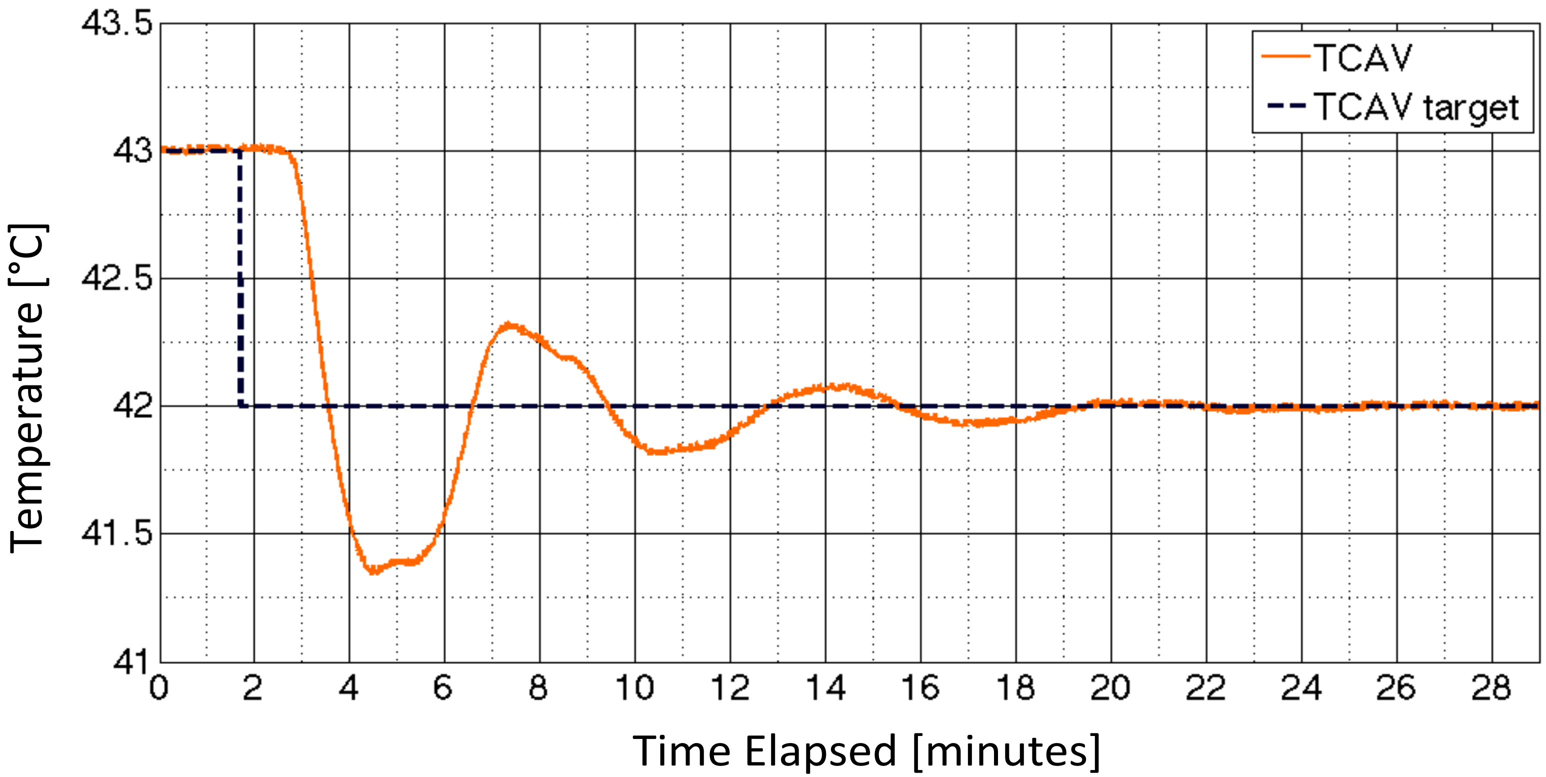

Note: oscillations are due to water recirculation + time delay (not PID tuning)

What can we do to improve this?

- Pre-emptive compensation for observed changes
- Account for time delays and use system history
- Use both the heater and the flow control valve

→ Model predictive control
→ Models identified from data

Model Predictive Control of the Water Temperature

supply temperature

water temperature exiting the gun

flow control valve settings

heater settings

temperature after mixing chamber (measured and target)

target cavity temperature
+ amount of RF power

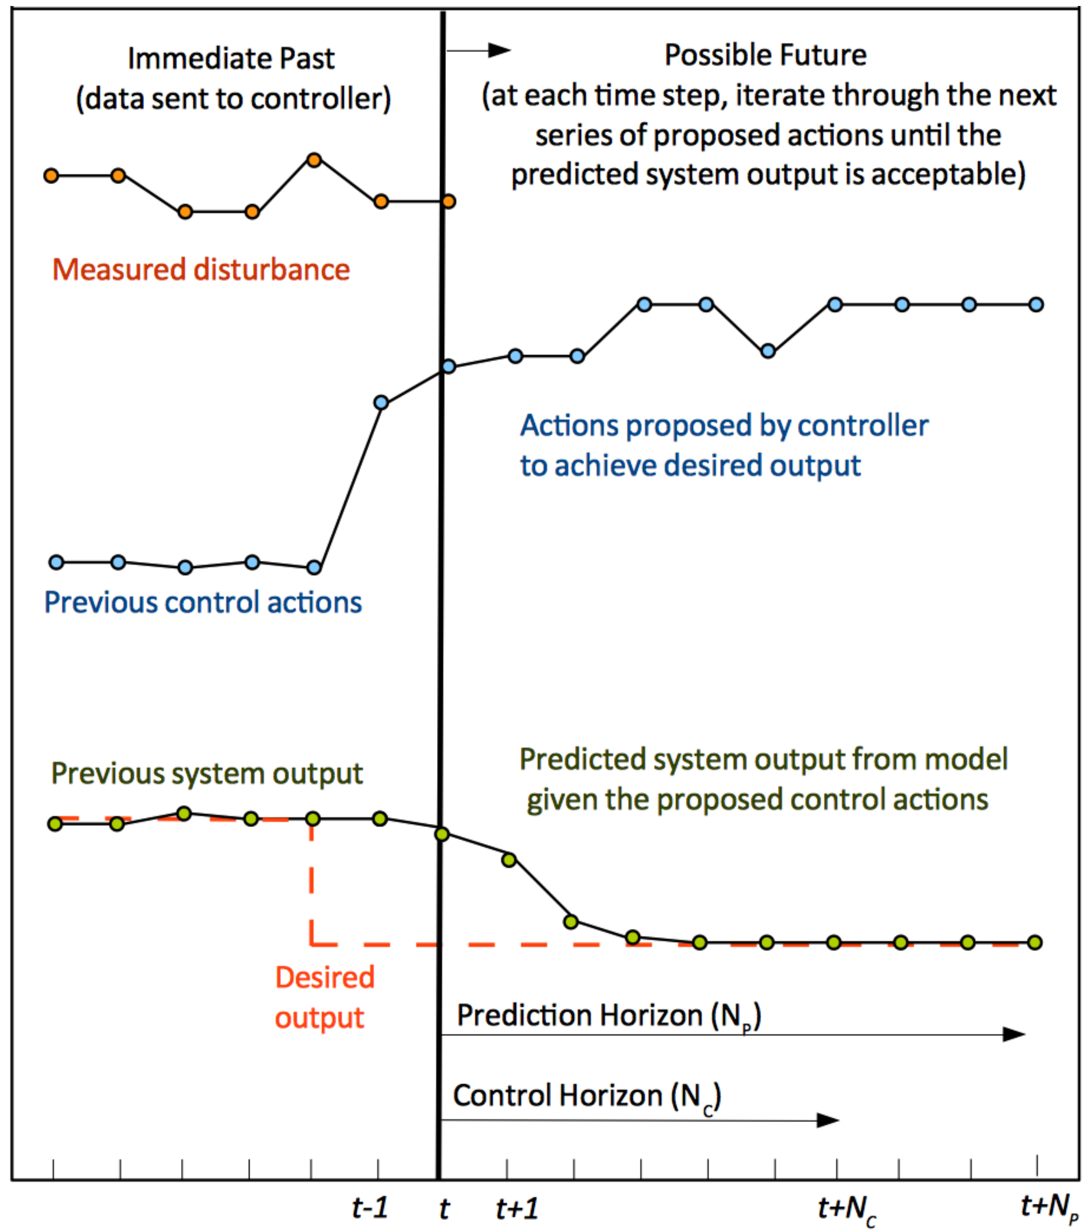

Input for Training and Validation Data

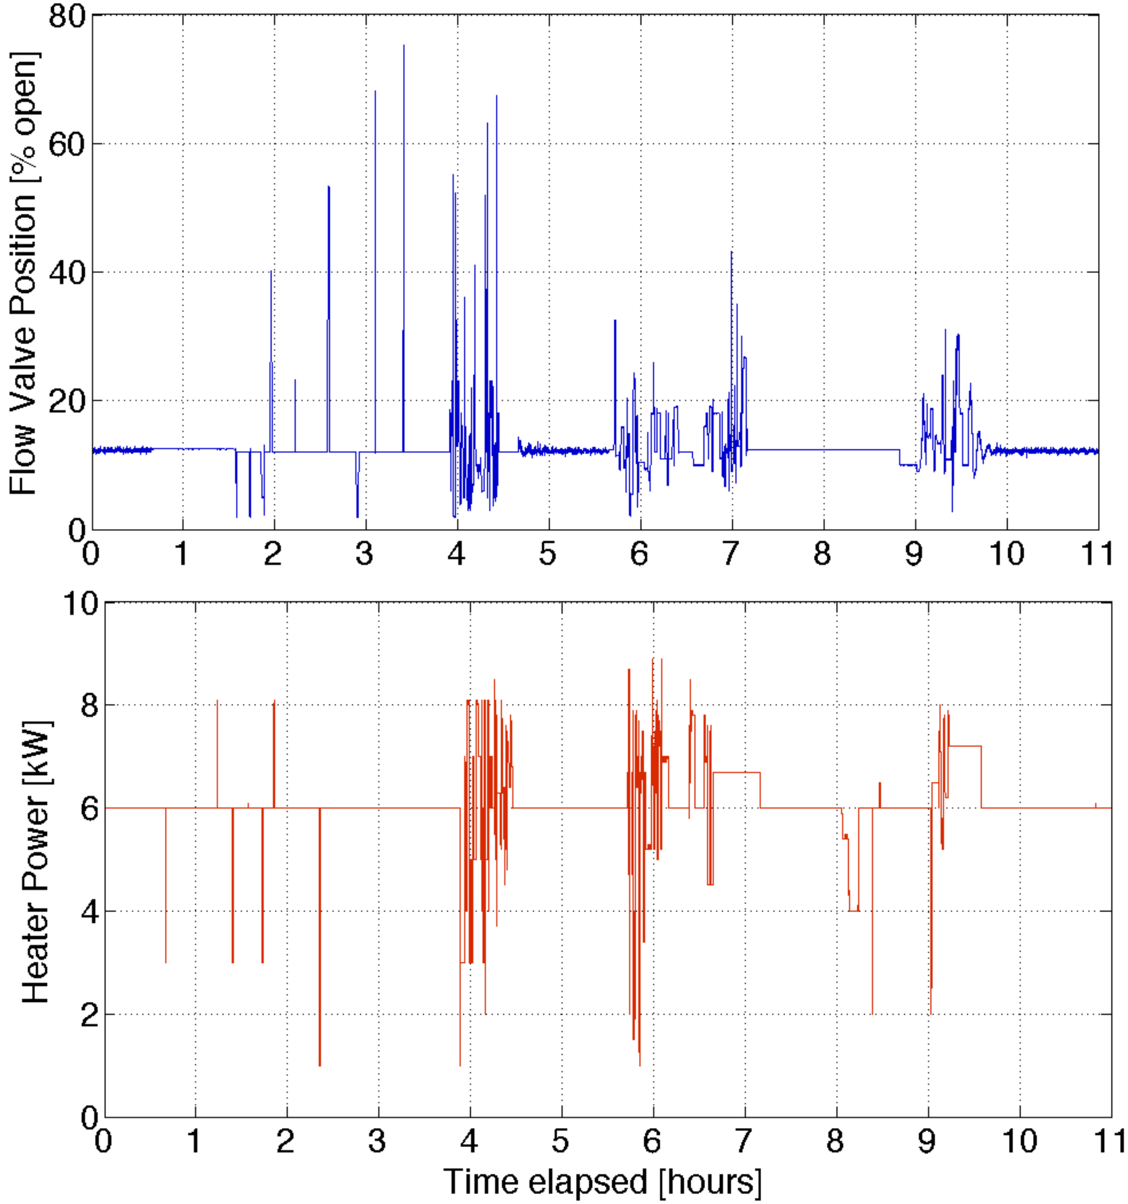

Note: there are more data sets than I am showing here

Model Performance (Training and Validation Data)

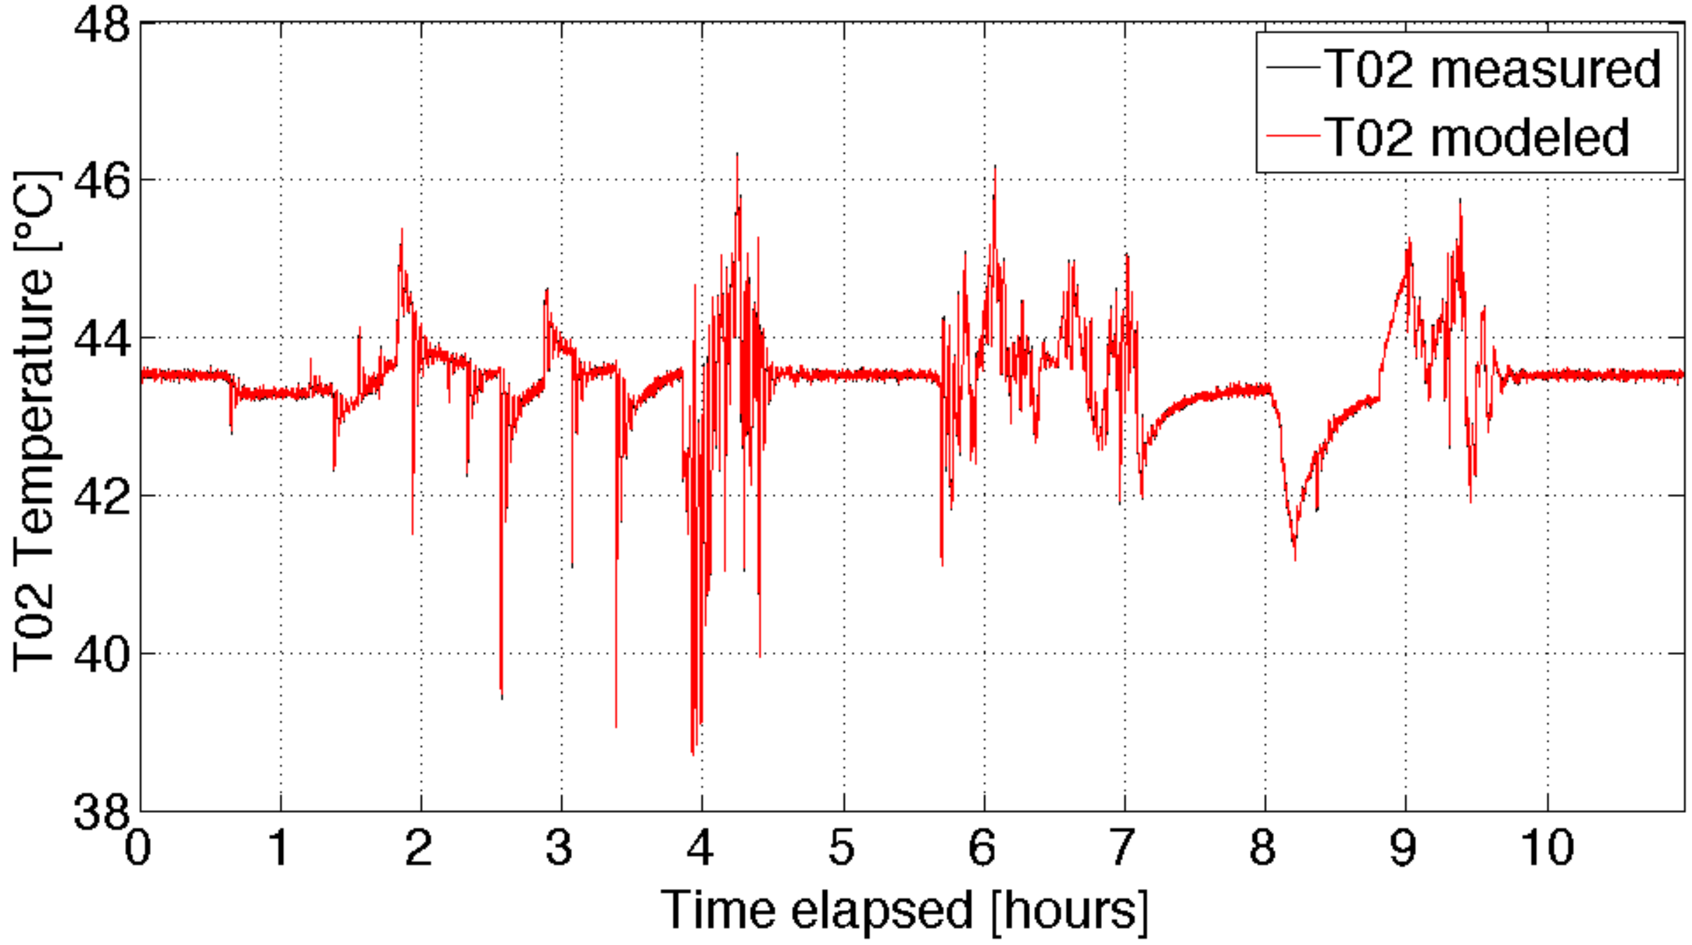

Testing Data: Change the RF Power Settings

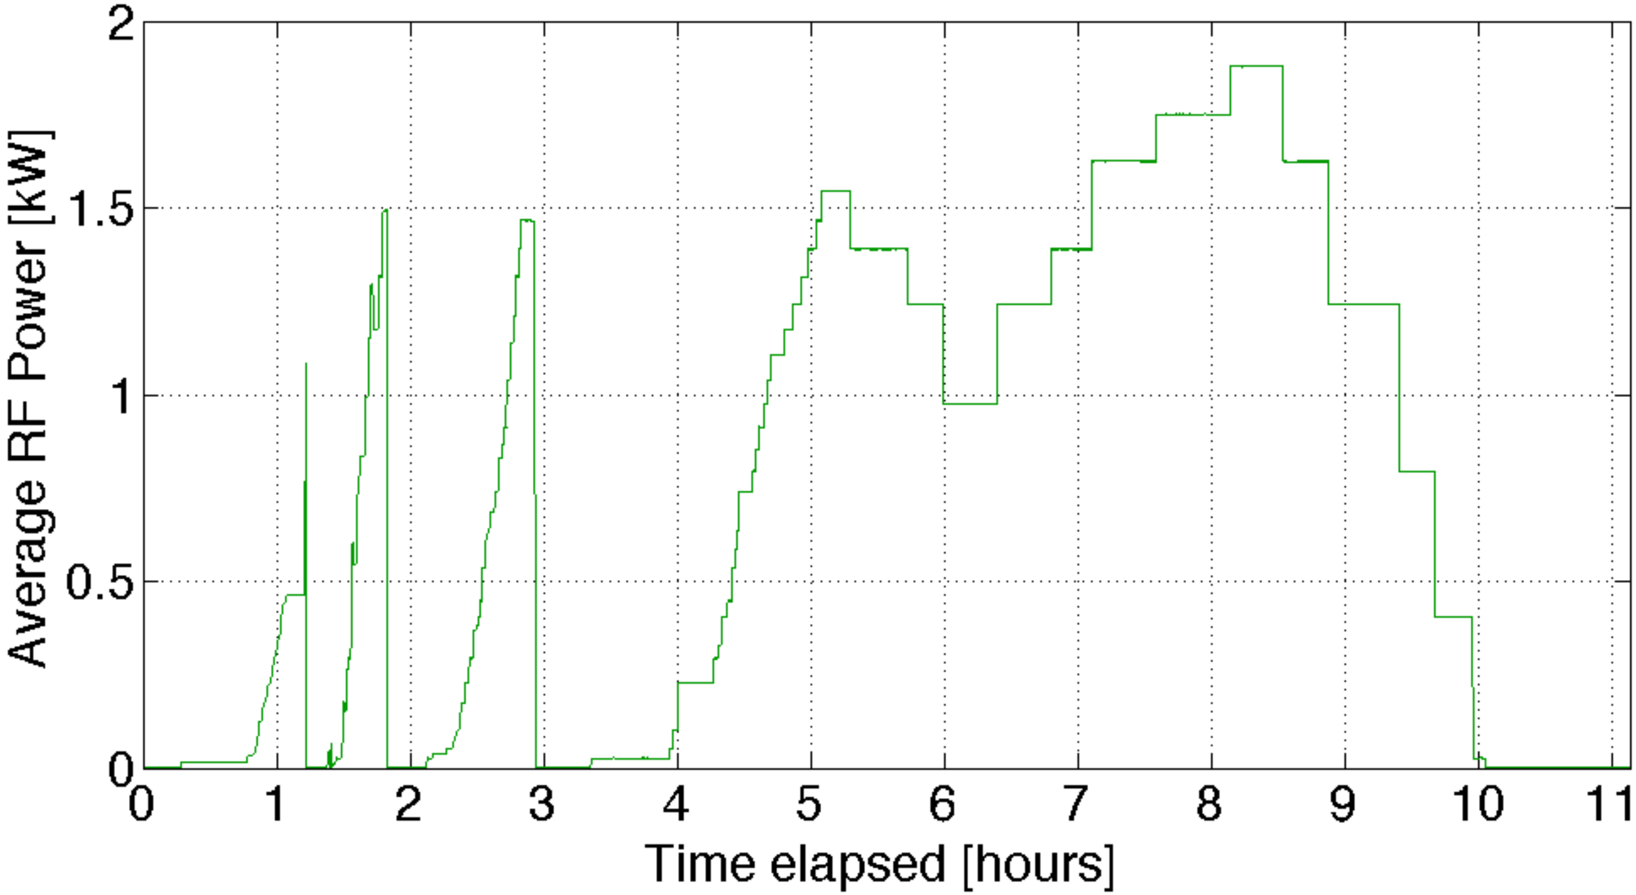

Model Performance (Testing Data)

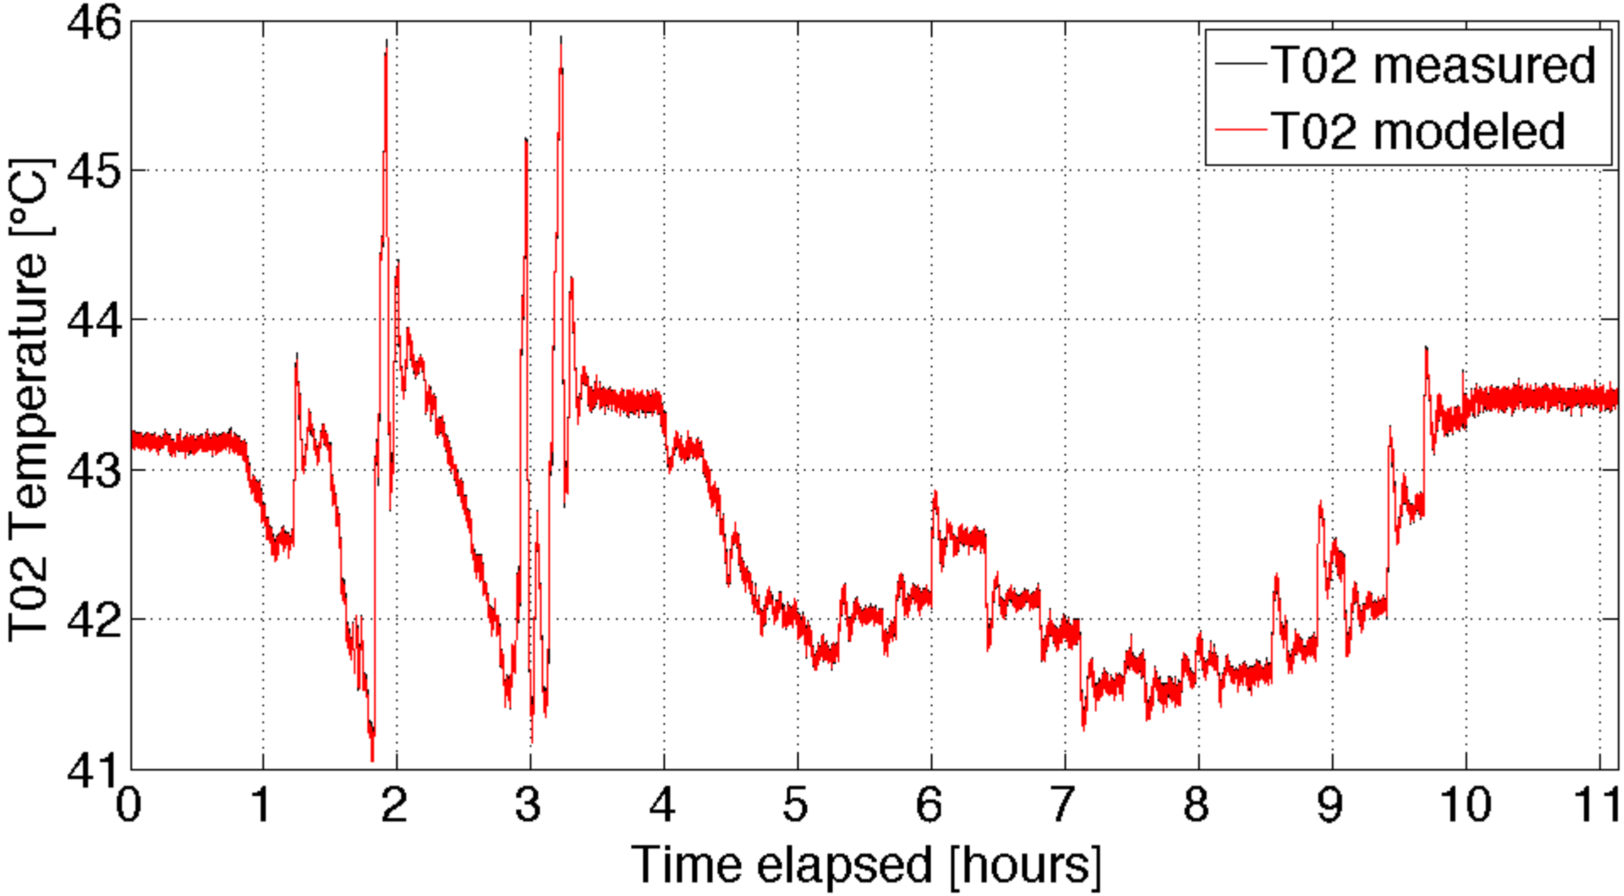

Preliminary MPC

- Performance benchmark to guide future design
- Reliable/fast optimization in a straightforward manner
  → simplified model of water temperature subsystem
- Rudimentary model for target water temperature based on average RF power and desired cavity temperature

Preliminary MPC: 1-°C Step Change

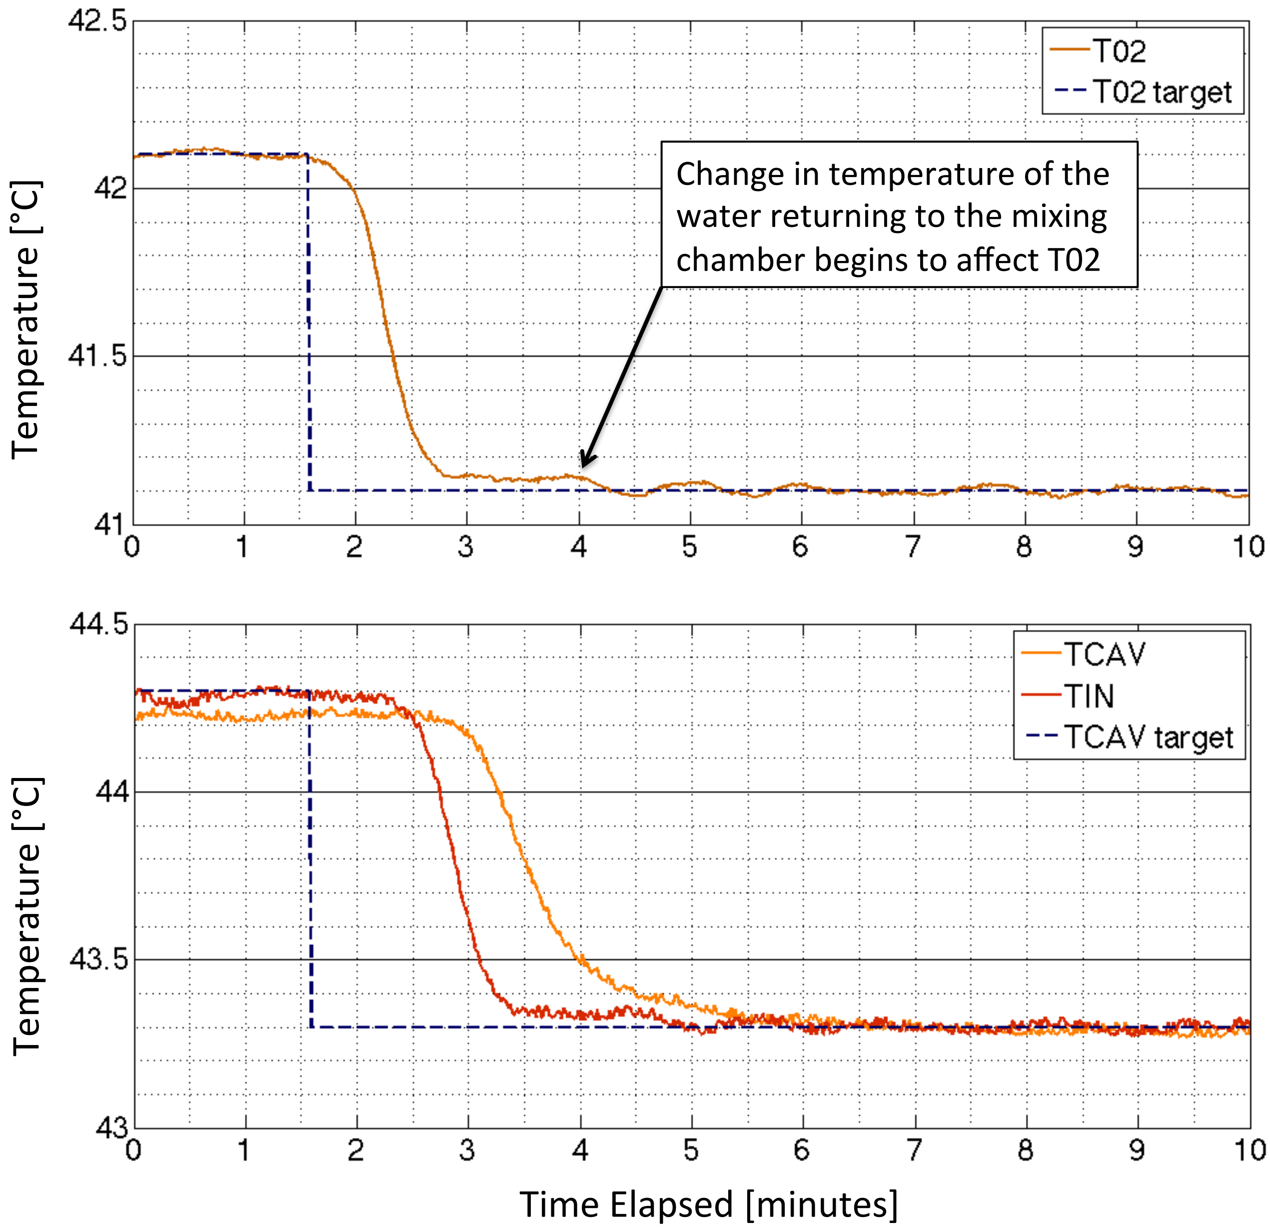

T02 within ±0.02 °C of its respective set point in about 3 minutes

TCAV within ±0.02 °C of its set point in about 5 minutes

~5x faster settling than PID

No large overshoot

Note:
-Difference in scale relative to the PID results

-There is some steady state offset in TCAV prior to the step

Preliminary MPC: 1-°C Step Change

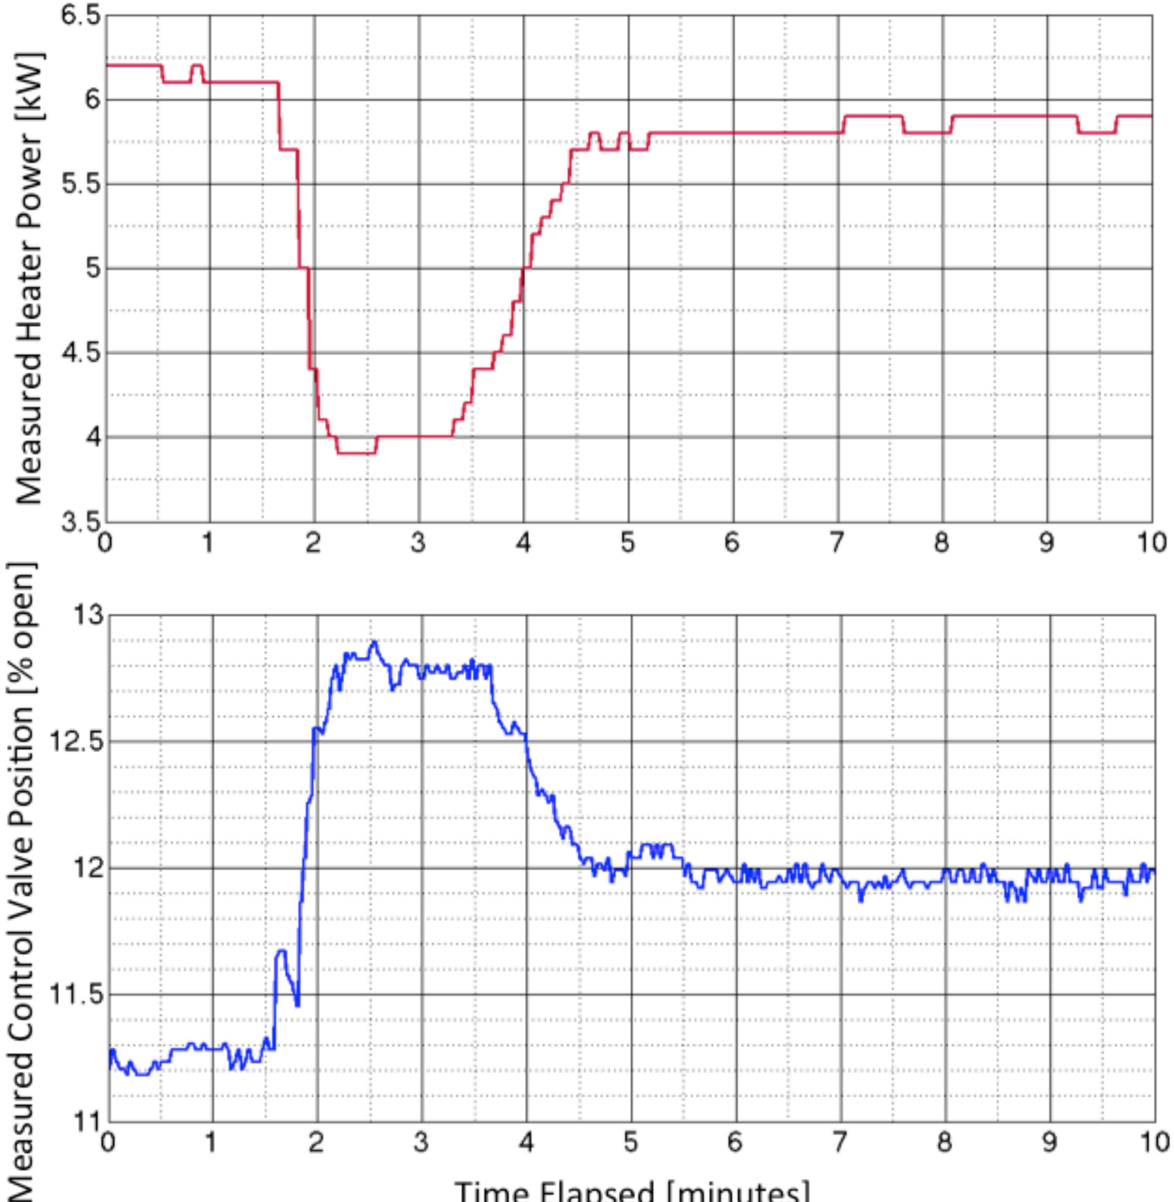

Next steps

- Improve the component that determines the water temperature set point
  - *account for variation in cave temperature*
- Clean up the implementation
  - *use measured actions, not just requested actions*
  - *online adaptation*
- Additional testing
  - *RF power*
  - *more complicated reference trajectory*
- If deemed necessary, use the more complicated/accurate subsystem model
- Develop resonance control component
  - *Forward and cavity phase measurements*
  - *Beam loading*
  - *Reflected power*
- Implementation for dedicated use